DOI: 10.1111/cgf.15200

COMPUTER GRAPHICS *forum*
Volume 43 (2024), number 6, e15200

EBPVis: Visual Analytics of Economic Behavior Patterns in a Virtual Experimental Environment

Yuhua Liu,[1] Yuming Ma,[1] Qing Shi,[2] Jin Wen,[1] Wanjun Zheng,[3] Xuanwu Yue,[4] Hang Ye,[3] Wei Chen,[5] Yuwei Meng[6] and Zhiguang Zhou[1,5]

[1]Intelligent Big Data Visualization Lab, Hangzhou Dianzi University, Hangzhou, China
{liuyuhua, mamingming, 231330023, zhgzhou}@hdu.edu.cn
[2]Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China
brantqshi@hkust-gz.edu.cn
[3]Zhejiang University of Finance & Economics, Hangzhou, China
zhengwanjun123@zufe.edu.cn, hangye@zju.edu.cn
[4]Sinovation Ventures AI Institute, Nanjing, China
xuanwu.yue@connect.ust.hk
[5]State Key Lab of CAD & CG, Zhejiang University, Hangzhou, China
chenvis@zju.edu.cn
[6]Zhejiang Provincial Energy Group Company, Hangzhou, China
mengyuwei@zhenergy.com.cn

Abstract

Experimental economics is an important branch of economics to study human behaviours in a controlled laboratory setting or out in the field. Scientific experiments are conducted in experimental economics to collect what decisions people make in specific circumstances and verify economic theories. As a significant couple of variables in the virtual experimental environment, decisions and outcomes change with the subjective factors of participants and objective circumstances, making it a difficult task to capture human behaviour patterns and establish correlations to verify economic theories. In this paper, we present a visual analytics system, EBPVis*, which enables economists to visually explore human behaviour patterns and faithfully verify economic theories,* e.g. *the vicious cycle of poverty and poverty trap. We utilize a Doc2Vec model to transform the economic behaviours of participants into a vectorized space according to their sequential decisions, where frequent sequences can be easily perceived and extracted to represent human behaviour patterns. To explore the correlation between decisions and outcomes, an Outcome View is designed to display the outcome variables for behaviour patterns. We also provide a Comparison View to support an efficient comparison between multiple behaviour patterns by revealing their differences in terms of decision combinations and time-varying profits. Moreover, an Individual View is designed to illustrate the outcome accumulation and behaviour patterns of subjects. Case studies, expert feedback and user studies based on a real-world dataset have demonstrated the effectiveness and practicability of* EBPVis *in the representation of economic behaviour patterns and certification of economic theories.*

Keywords: visualization; visual analytics, information visualization

CCS Concepts: • Human-centred computing → Visual analytics; Information visualization

1. Introduction

As an important branch of economics, experimental economics investigates individuals' economic behaviours under controlled conditions in the laboratory or out in the field. It simulates the real-world environment to observe and analyse decisions made by subjects motivated by practical rewards for verifying economic theories. Currently, experimental economics has been used to address economic issues such as the market mechanism, poverty trap and salary administration [Org15, TT12]. To enhance the experiments' validity [Bol10], economists often implement a set of scenarios and real-effort tasks [CD17] to simulate real-life situations [BS07].

We collaborate closely with experimental economists to engage in a laboratory experiment, focusing on investigating risk decision-making and preferences of individuals under different rich and poor

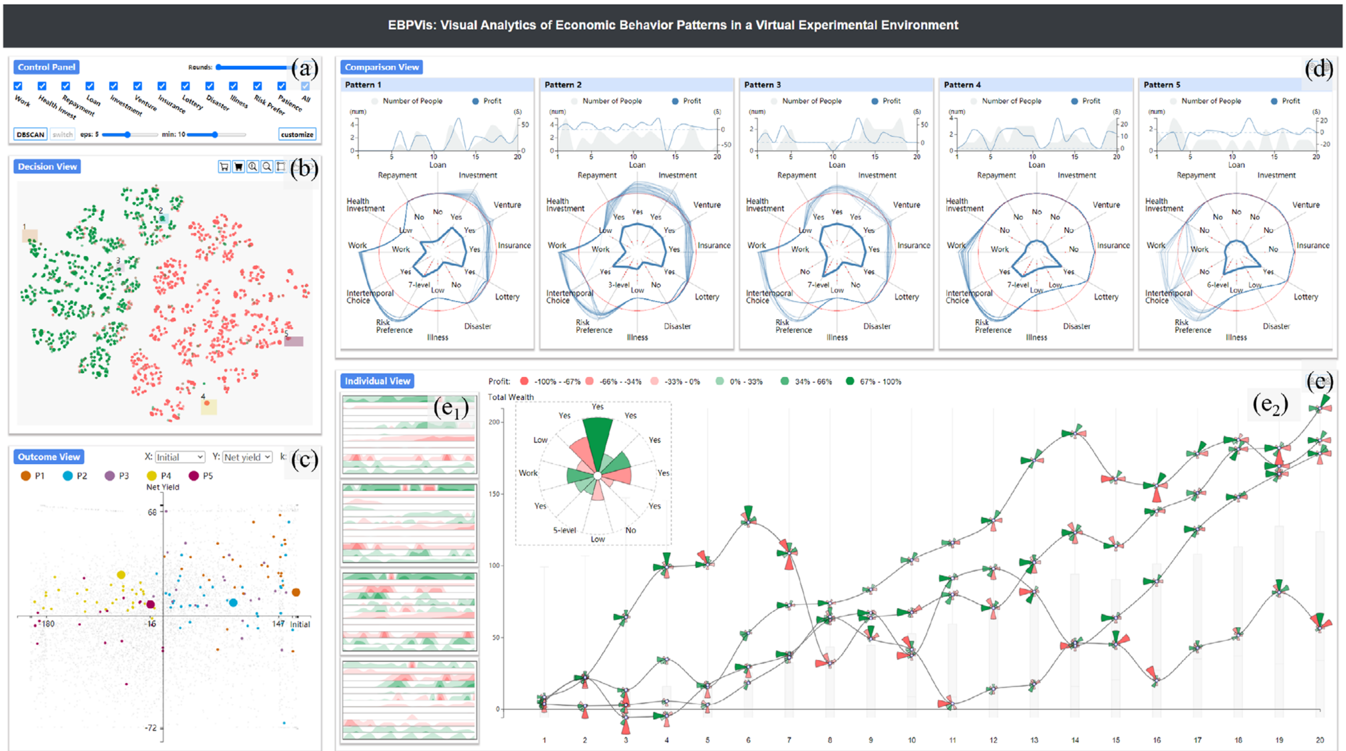

Figure 1: *A case study for visual analytics of economic behaviour patterns in a laboratory experiment to investigate the risk decision-making and preferences of individuals in rich and poor states. A behaviour pattern extraction model is conducted in the control panel (a), to transform the economic behaviours into an embedding space (b), enabling users to easily perceive and extract behaviour patterns. The Outcome View (c) reveals the correlation between decision behaviours and wealth, enabling users to focus on patterns of interest. Behaviour patterns are visually compared in the Comparison View (d) in terms of decision combination and profit. The Individual View (e) illustrates the wealth accumulation and behaviour patterns of subjects of interest, to further verify economic theories.*

states. A series of real-life scenarios, such as investment, insurance, loan and work, are incorporated into the experiment and played repetitively over dozens of rounds. Each subject makes a choice (*i.e.* decision), among multiple alternatives in each situation. These decisions all affect the subjects' wealth accumulations (*i.e.* outcomes), such as buying lottery tickets, seeking medical attention and purchasing insurance. The experimental procedures involve a substantial number of experimental variables. What's more, there are often hundreds of subjects recruited to participate in the experiment, who behave diversely and unpredictably. Therefore, the experimental data are collected with the characteristics of high-dimension, time-varying and large uncertainty, which makes it a difficult and time-consuming task to capture human behaviour patterns and establish correlations to verify economic theories, *e.g.* the vicious cycle of poverty [MV05] and poverty trap [DC08]. According to in-depth discussions with domain experts, the challenges for the exploration of experimental economics data mainly come from three aspects as follows:

C1: Dynamics and diversity of decision-making behaviours. With cognitive biases, personal preferences and external situations, subjects always present various behaviours during the same round. For each subject, the decisions he/she made might be different across repeated sessions [PWLM22]. Thus, the decision-making behaviours are highly dynamic and diversified, which brings great challenges for the extraction and analysis of behaviour patterns.

C2: Mutual effect between decision-making behaviours and outcomes. Subjects often make decisions according to their economic conditions. Meanwhile, the decisions made by the subjects also have effects on the outcome accumulations. For example, a series of decisions made by a subject in different sessions would lead to changes in his outcome. At the same time, different levels of outcomes would also guide the subject to change decisions to maximize his outcome. Such mutual correlations are complicated to model for further certification of economic theories.

C3: Limitation of traditional analytical methods. Economists typically explore experimental data using statistical models and software like SPSS [SPS90] and R [DP08]. Such methods either present limited results without suitable interpretations or require experienced programming skills for desired visualizations and analysis, which brings much trouble for experts to interactively explore the dynamic processes of decision-making, decision adjustments, and outcome changes according to their domain knowledge.

To address the above challenges, we develop a visual analytics system, *EBPVis* (Figure 1), enabling economists to explore the behaviour patterns of subjects contained in the experimental economics data. First, we utilize a Doc2Vec [LM14] model to transform subjects' behaviours into a low-dimensional embedding space according to their sequential decisions, in which the underlying behaviour patterns can be visually captured and easily extracted **(C1)**. Second, a coordinates-based Outcome View is designed to measure the wealth variables of behaviour patterns, allowing the economists to visually explore the correlation between human behaviour patterns and outcomes **(C2)**. Moreover, a Comparison View with nested radar maps is proposed to quickly identify the differences between multiple behaviour patterns, which provides a comprehensive and intuitive comparison of behaviour patterns from aspects of decision combinations and time-varying profits **(C2)**. Meanwhile, an Individual View comprised of glyph-embedded line charts and horizon graphs is designed to illustrate the evolutionary behaviours and outcomes of individuals in detail **(C3)**. Case studies based on a real-world dataset have demonstrated the validity of our system in visually exploring human behaviour patterns and faithfully verifying economic theories from multiple perspectives. In summary, the major contributions of this paper are listed as follows:

- To the best of our knowledge, we are the first to design visual analytics technologies for exploring experimental economics data and verifying economic theories (*i.e*. poverty trap).
- A deep text representation model (*i.e*. Doc2Vec) is utilized to capture the economic behaviour patterns according to the co-occurrence likelihood of decisions and outcomes.
- A visual analytics system, *EBPVis*, integrates effective visualizations and interactions designed to facilitate the establishment of correlations between economic behaviours and outcomes.

2. Related Work

2.1. Economic behaviour analysis

Empirical analysis in behavioural economics tests assumptions by observing actual behaviours [WBK19]. As a variance, experimental economics aims to study people's economic behaviour in a controlled laboratory setting or out in the field. For example, Bo *et al.* [BQL*20] explored managers' behaviour in a controlled economic experiment affecting the distribution mechanism of medical delivery systems. Tian *et al.* [TCS19] conducted an online interactive experiment to study the behavioural effects of a tradable mobility credit scheme. Recently, visualization techniques have been widely applied to analyse economic behaviours [KCA*16]. For example, Gunaratne and Nov [GN15] developed a user interface for retirement decision-making based on endowment effect and loss aversion to help users adjust their behaviours. Wesslen *et al.* [WKMD21] conduct a controlled and incentivized crowdsourced experiment replicating Benartzi and Thaler [TB99] and extending it by testing the effect of a range of new uncertainty visualizations on myopic loss aversion. Savikhin *et al.* [SLFE11] exploited Portfolio-Compare, an interactive visual analytic decision support tool for analysing risk and return in economic experiments. Csallner *et al.* [CHLS03] developed FundExplorer, using context treemaps to enhance the transparency of equity mutual fund portfolios and aid diversification. To help non-expert users interpret financial data, Rudolph *et al.* [RSE09] proposed a visual analytics tool, FinVis, through which the decision-maker quickly chooses between financial portfolio options. To the best of our knowledge, there are few studies focusing on the visual exploration and analysis of experimental economics data in the visualization community. This paper implements a visualization system to explore the economic behaviour patterns located in the experimental economics data.

2.2. Event sequence visualization

Event sequence data refers to a series of discrete events, which are temporally ordered according to their occurrence. A variety of visualization methods have been proposed to explore the patterns of event sequence data [ZYC*24]. Guo *et al.* [GGJ*21] reviewed the visual analytics methods for event sequence data exploration, and classified them into representation, interaction technology and application. Chen *et al.* [CXR17] proposed a visual summary of event sequence data with the minimum description length principle, enabling users' detailed exploration. Guo *et al.* [GXZ*17] introduced EventThread, a visualization system specializing in healthcare, offering visual summarization and stage analysis of event sequence data. They also designed an unsupervised stage analysis algorithm [GJG*18] to uncover key stages in event sequences, using visualization to reveal evolutionary patterns across these stages. Gotz *et al.* [GZW*19] proposed a visualization method for the hierarchical exploration of high-dimensional medical event sequence data. Cappers *et al.* [CvW17] presented an approach to facilitate both temporal and structural analysis of multivariate event sequence data. Liu *et al.* [LWD*16] proposed an analysis pipeline to mine, prune and explore clickstream data to understand common visitor paths. Law *et al.* [LLMB18] developed a mining and querying user interface, MAQUI, allowing the analysts to recursively explore event sequence data. In contrast, our work not only considers the sequential features in the decision behaviours, but also takes into account the associated wealth features before and after each decision toward a better comprehension of economic behaviour patterns.

2.3. Association rule mining

Association rule mining (ARM) is a most popular research topic in the field of big data analysis, the goal of which is to discover association rules in the transaction database. It has many real-life applications, such as marketing strategies, cause analysis of accidents and news recommendations. Based on the well-known Apriori algorithm [AS*94], a variety of ARM algorithms have been proposed to improve the association rule interestingness measures [LVML07] and performance [HPY00]. Two phases are often included like the mining of frequent item sets and the generation of rules with these item sets. However, the occurrence time or sequential ordering of events is ignored in the process. Thus, sequential rule mining (SRM) was proposed, which focused on the exploration of sequential rules according to sequential ordering, aiming at the solution of prediction-related tasks. Many SRM algorithms have been proposed to mine the sequential rules by means of a pattern-growth or a vertical approach. Based on the SRM, a number of related topics are studied such as frequent sequential patterns, maximal sequential patterns and closed sequential patterns [FVLK*17]. In this paper, the association between decision behaviours and outcomes in economic experiments is also a focus of our work.

2.4. Visual analysis of multi-dimensional data

Multi-dimensional datasets are widely collected in a variety of application domains, such as biology, economics, medicine, statistics and chemistry. Various visualization methods have been implemented to convey the multi-dimensional structural information [LMW*16, MKW24]. In these methods, scatterplot matrix [KHD*10, HSvK*19a] and parallel coordinate plots [RLS*18, ZMZ*21] are widely used to discover the features and trends of multi-dimensional data. The scatterplot matrix is a collection of bivariate scatterplots, allowing users to easily analyse the relationship between any pair of dimensions easily. In contrast, parallel coordinate plots display all axes simultaneously, with each data item represented as a polygonal line connecting corresponding values on the axes, better revealing multivariate relationships. Other methods can also be applied to explore the relationship pattern for multi-dimensional data, such as star coordinates [Kan00], Radviz [HGM*97] and glyphs [BKC*13]. To reduce the visual clutter of large-scale multi-dimensional data, various extensions are developed to enhance multi-dimensional data analysis based on the axes recording [HW13], edge bundling [ZYQ*08], sampling [HSVK*19b], clustering [LWCC17], *etc*. The decisions and outcomes in the economic experiments are characterized as multi-dimensional time-series data. Effectively displaying and comparing them plays a crucial role in exploring experimental economic data.

3. Requirement Analysis and System Overview

3.1. Background

We worked closely with the director of the Center for Economic Behavior and Decision-making[1] at one university. The director and his team have long been dedicated to the interdisciplinary study of economics, spanning behavioural economics, experimental economics and computer simulation of human social behaviours. Recently, they conducted an economic experiment exploring the endogenous process of biochemical poverty to investigate risk preferences of individuals in different rich and poor states. They have recruited 304 subjects (160 males and 144 females) from different grades (including freshmen to seniors) and different majors (including economics, law, computer science, *etc*.). Each subject is provided with a computer equipped with the experimental system, and will be introduced with the use of the system.

As depicted in Figure 2, this experiment consisted of a series of scenario-based simulations conducted over 20 rounds. In each round, there were 13 sequential sessions covering scenarios like work and health investment. Among them, the work session introduced a real-effort task where subjects earned wealth by answering questions. To make the experiment more realistic, subjects had to consume constant wealth on living expenses to survive in the work session. An unemployment session was designed so that each subject had a certain probability of being unemployed. In the investment session, subjects could invest their own money, and the rewards of their investments were decided by a probabilistic model. The disaster session simulated the loss of property caused by disasters such

[1]https://cebd.zufe.edu.cn/English/Home.htm

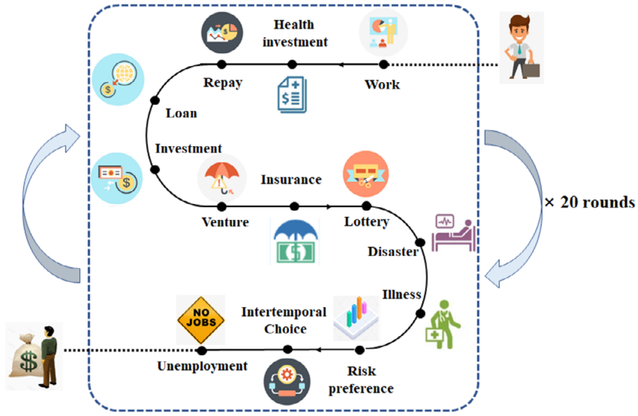

Figure 2: *Experimental procedures are designed to study the correlation between human behaviour patterns and outcomes.*

as earthquakes, floods, typhoons and fires in life. In the session on risk preference, the subjects had to undergo a process of testing to determine their risk preferences. The experimental system offered subjects various risk-level choices, with outcomes determined by coin tosses. Other sessions aimed to simulate real-life situations and provide a good sense of substitution and reality as far as possible. Alternative choices in each session and the corresponding profits are summarized in the supplementary material.

It should be noted that the round number of the experiment was formulated by our cooperative experts. The experts were interviewed and told us that they had carried out a lot of explorations about the round number of the experiment. Finally, it was found that 20 rounds could just widen the subjects' wealth gap effectively, which was sufficient for them to investigate the risk decision-making and preferences of individuals under different rich and poor states. With fewer rounds, the wealth gap is not obvious. The subjects have to spend more time on repeated experiments with too many rounds. Statistical results have shown that it took about 6 min to complete a round, and each subject spent about 2 h on the whole experiment. Another point is that the order of the 13 experimental sessions was fixed by domain experts. The former sessions such as work, investment, and insurance were to create wealth and protect wealth, while the latter sessions such as disaster and illness were to consume wealth. Such a fixed-sequence experiment could reduce the complexity of scenarios, making it easier for subjects to understand the rules and get familiar with the experiment. In addition, the initial wealth of each subject was 0 before the first round. The wealth earned by the subjects was not cleared after each round, and it would be carried to the next round for investment, insurance and other sessions. All subjects were compensated for their participation in this study. Particularly, the subjects who gained virtual wealth in the experiment would receive real money. The more wealth one subject earned, the more reward he/she received. The motive to earn money will enable subjects to make efforts to earn more wealth in the experiment. Finally, the economic behaviour data is collected by the experimental system, which contains 304 (subjects) × 20 (rounds) = 6080 records. Each record holds the decision information and the associated wealth information of a subject in a round of experiments.

3.2. Requirement analysis

Given the experimental data, the domain experts (from the Center for Economic Behavior and Decision-making) always utilize traditional data analysis tools like SPSS [SPS90], STATA [CT*10] and R [DP08] to focus on some fragments such as the session of risk preference in each round (Figure 2), and study the risk preference of individuals in different rich and poor states. However, it is still a difficult task to take advantage of both the multi-dimensional decisions and time-varying outcomes in the experimental data, get deep insights into the economic behaviours and draw comprehensive conclusions from the economic theories. Therefore, the experts claimed that it would be wonderful to design a visual analytics tool to explore intricate decision behaviours in addition to their effects on the decision sequence and wealth. This would yield valuable insights to improve experimental parameters and verify poverty-related economic theories like the poverty trap and the vicious circle of poverty. We have held seminars with the experts twice a month in the past half year. The entire process was divided into two phases, each consisting of extensive formal interviews with all or partial experts. The first month was the first phase, which was to identify the expert needs and analysis tasks in their exploration of economic behaviour patterns. The following 5 months involved iterative development of the visual analytics system based on expert feedback. Four requirements were summarized based on our discussion with the experts, which are listed as follows:

R1. How to extract the behaviour patterns? In each experiment round, the behaviour of each subject is characterized by a series of decisions on multiple sessions. The decision combinations covering all subjects and all rounds make it complicated to extract the behaviours of subjects. In addition, the associated wealth also presents the underlying nature of behaviours. Thus, identifying behaviour patterns with both decisions and wealth considered is crucial for economists to reduce their cognitive burden of massive decision behaviours.

R2. How to display and compare the behaviour patterns? It is always found that some behaviours share similar patterns with each other in terms of decisions and wealth effect. Thus, it is significant to provide economists insights into the dimensions of decisions and wealth, for the enhancement of the awareness of behaviour patterns. Besides, revealing the behaviour pattern from aspects of distribution characteristics, temporal trends and wealth profits will also facilitate economists to compare and explore the inter-pattern difference.

R3. What is the correlation between behaviour patterns and wealth? The behaviours of the subjects in the experiment exhibit a close correlation with the changes in wealth. Which behaviour pattern can help the subjects to gain profits as far as possible? How are the behaviours limited by the individual wealth condition? Is there any difference in behaviours between the poor and the rich states? It is an important task to understand the rationality behind the behaviour patterns, enabling economists to analyse the intricate correlations between behaviours and wealth.

R4. How do the behaviour patterns and wealth of a subject evolve over time? Behaviour patterns of the subjects vary with time according to their different decisions on specific experimental sessions. Which decision mainly affects the wealth profit in a round of experiments? How are the behaviour patterns of the subjects transformed from one to another? It is an inevitable task to display such individual-level information in detail, enabling economists to get deeper insights into common and unique behaviour patterns.

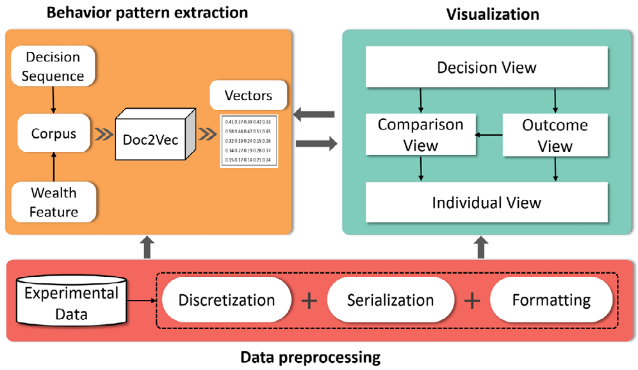

Figure 3: *The system pipeline of* EBPVis*. Three modules are included: (1) Raw experimental data are processed in the data preprocessing module. (2) The behaviour pattern extraction module represents the behaviours as vectors. (3) The visualization module is provided to explore behaviour patterns.*

3.3. System overview

The system pipeline of *EBPVis* is presented in Figure 3. Three components are included: data pre-processing module, behaviour pattern extraction module and visualization module. The raw experimental data are loaded and processed in the data pre-processing module. The behaviour pattern extraction module is implemented based on a free Python library Gensim, which includes a set of Natural Language Processing (NLP) models. In particular, the Doc2Vec model is applied to manipulate the processed data for the extraction of high-level behaviour patterns. The visualization interface is composed of four well-structured visualization views, enabling users to explore the behaviour patterns and the associated factors fully.

4. Economic Behaviour Pattern Extraction

According to the above discussions, experimental data have high-dimensional and time-varying characteristics. Traditional sequential pattern mining (SRM) methods [FVLK*17] focus on identifying the longest and most frequent patterns in an automated way. Applying these methods to extract behaviour patterns still suffers from several problems. One is that the relation between the identified patterns is difficult to quantify and convey. Economists need to explore the behaviour patterns considering all the decisions in the context of the entire experiment. Another point is that SRM is sensitive to parameters such as minimum support and minimum length, which often requires users to undergo a trial-and-error process, thus forming better quality patterns.

In the field of NLP, embedding techniques have been widely used to analyse and model the relational context in text, such as Word2Vec [MCCD13] and Doc2Vec [LM14]. Word2Vec learns distributed representations of words following the rationale that words

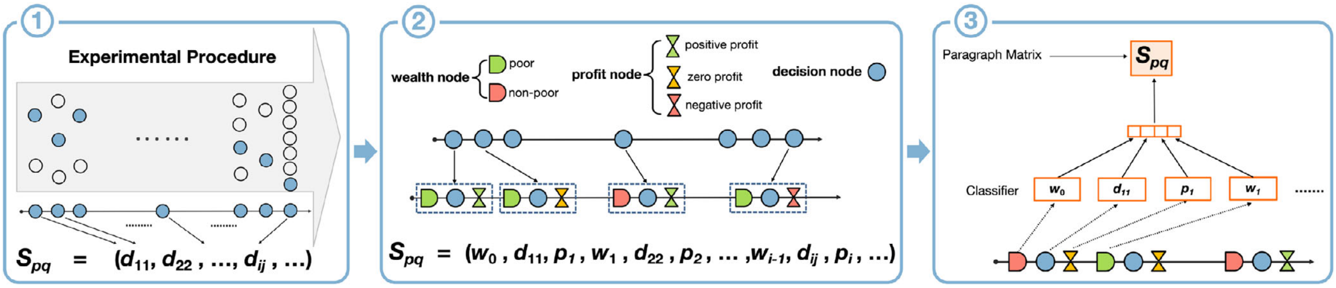

Figure 4: *The illustration of behaviour pattern extraction: (1) pure decision sequence, (2) wealth-enhanced decision sequence integrating the decision and wealth information and (3) sequence representation.*

sharing similar contexts tend to have similar vector representations. As an extension of Word2Vec, Doc2Vec incorporates document or sentence IDs into the word sequences to maximize their co-occurrence likelihood with the representations of documents and words learned simultaneously. Due to fully exploiting the semantics of words in documents, Doc2Vec can map documents with similar contextual structures to a similar position in the vector space. In addition to representing texts, Doc2Vec is capable of constructing representations for sequential data, just as the economic behaviours studied in this paper. By contrast, the standard econometric models are mainly linear regression, binary logit and multinomial (or nested) logit, which focuses on identifying causal effects, testing hypotheses and estimating parameters [DE20] [ABRS21]. They are incapable of semantic description and semantic analysis for sequence data, thus the common characteristics in the numerous and complicated economic behaviours cannot be effectively mined and explored. Therefore, Doc2Vec is leveraged here to discover and extract underlying behaviour patterns combining both sequential features from the decisions as well as the associated wealth features.

4.1. Corpus generation

The raw experimental data record the detailed decisions covering 304 subjects and 20 rounds. The behaviour made by the pth subject during the qth experimental round is denoted as S_{pq} ($1 \leq p \leq 304$, $1 \leq q \leq 20$). As shown in Figure 2, the experiment designs 13 sessions in each round, ranging from work to inter-temporal choice and risk preference. In particular, the session of unemployment determines whether a subject can work in the next round. So we remove the unemployment session from the current round, and combine it with the work session of the next round into a single session, which consists of two choices: employment and unemployment. Finally, 12 sessions in total are reserved to construct the decision sequence, which is numbered $i = 1, 2, \ldots, 12$. Therefore the behaviour of one subject in a round can be characterized by a sequence of decisions. As shown in Figure 4(1), $S_{pq} = (d_{11}, d_{22}, \ldots, d_{ij}, \ldots)$ and d_{ij} indicates that the pth subject selects the jth choice in the ith session.

However, decision features are still insufficient for the exploration of economic behaviours and their correlations with outcomes. The identical behaviours corresponding to the same decision sequence would also produce significantly different outcomes. For example, two subjects both choose 'Yes' in the session of investment, but their outcomes are randomly set by the experimental system. The lucky subject might receive a positive outcome while the other with bad luck would lose wealth. Similar things often occur in the other sessions, which have a variety of impacts on the final outcome. Our collaborative experts wanted us to distinguish between such behaviours by precisely describing and analysing the derived patterns in terms of both the decision features and wealth features.

To meet this requirement, we define another two types of nodes in addition to the decision node, and incorporate them into the pure decision sequence, as illustrated in Figure 4(2). Before each session, we classify the wealth of subjects into two categories: poor and non-poor. Similarly, the profits after each session are grouped into three categories: positive profit, zero profit and negative profit. Each decision node d_{ij} is, respectively, linked with a wealth node w_{i-1} ahead of it and a profit node p_i behind it, representing what kind of wealth condition the subject possesses before he makes a choice in the session and whether he earns some money after he finishes the session. The behaviour made by the pth subject in the qth round is initially composed of a series of pure decision nodes. After the wealth nodes and profit nodes are inserted into it (Figure 4(2)), S_{pq} is revised as $\{w_0, d_{11}, p_1, w_1, d_{22}, p_2, w_2, \ldots, w_{i-1}, d_{ij}, p_i \ldots\}$ and $w_0, w_1, \ldots, w_i, \ldots \in \{\text{poor, non-poor}\}, p_1, p_2, \ldots, p_i, \ldots, \{$positive-profit, zero-profit and negative-profit$\}$, which contains complementary information about the initial wealth and profit. The behaviours of all subjects across all rounds are transformed into such wealth-enhanced decision sequences by the way described above, which generate a large and complete corpus.

4.2. Doc2Vec-based pattern extraction

In our behaviour pattern extraction model, we regard behaviours as documents structured by decision nodes, wealth nodes and profit nodes which, in turn, are analogical words from a special language. Each behaviour $S_{pq} = \{w_0, d_{11}, p_1, w_1, d_{22}, p_2, w_2, \ldots, w_{i-1}, d_{ij}, p_i, \ldots\}$ is modelled as a document $doc_m = \{word_1, word_2, \ldots, word_n, \ldots\}$. The skipgram model, paragraph vector-distributed bag of words (PV-DM), is then applied to learn behaviour embeddings. As shown in Figure 4(3), treating $w_n \in d_m$ as occurring in the context of d_m, we aim to maximize the following log-likelihood:

$$\frac{1}{N} \sum_{i=k}^{N-k} \log p(w_i | d_i, w_{i-k}, \ldots, w_{i+k}) \tag{1}$$

where N stands for the length of the behaviour sequence, k is the size of the text window and the prediction task can be completed using a softmax classifier:

$$\log p(w_i|d_i, w_{i-k}, \ldots, w_{i+k}) = \frac{e^{Z_{w_i}}}{\sum_j e^{Z_j}} \quad (2)$$

$$Z = b + Uh(d_i, w_{i-k}, \ldots, w_{i+k}; W, D) \quad (3)$$

Each of Z_j is un-normalized log-probability for each output word j, where U and b are the basic parameters of the softmax and h is derived by averaging or concatenating the word matrix W and paragraph vector D. After training using stochastic gradient descent and backpropagation, the document vectors, namely the behaviour vectors, are learned. These vectors, inferred by fixing features of decisions, profits and wealth conditions, map similar behaviours to similar positions in the vector space, which indicates the subjects make consistent decisions and gain or lose wealth synchronously on multiple sessions in the same rich or poor states. We compared Doc2Vec with traditional SRM in behaviour capture, as detailed in the supplementary material. The results indicate that Doc2Vec can more effectively learn and quantify contextual features of decision sequences.

4.3. Experimental comparison

To further prove the necessity of considering wealth information, we conducted experiments comparing clustering results with and without considering the wealth and profit information, which proves that introducing the wealth and profit information into the data sequence will generate better clusters. Two types of data sequences, *i.e.* whether adding the wealth nodes and profit nodes into the decision sequences, are, respectively, trained by the Doc2Vec model and transformed into vectors. Then the K-Means method is utilized to group the vectors into different clusters. To quantify the homogeneity of wealth, decision and profit in each cluster, an evaluation metric, average clustering entropy, is defined as follows:

$$entropy = -\frac{1}{S}\sum_{i=1}^{S}\sum_{j=1}^{|A_i|} p_{ij}\log p_{ij} \quad (4)$$

where S is the number of experimental sessions and $A \in$ {wealth, decision, profit}. In addition, p_{ij} denotes the percentage of behaviours in this cluster whose wealth/decision/profit category on the ith $session = a_{ij}$, where $a_{ij} \in Dom(Ai)$. Thus, this metric represents the average value of information entropies covering all sessions with respect to wealth features, behaviour decisions or profit levels. A small entropy value indicates a better clustering result.

Figure 5 gives an overview of the scores of clusters under different numbers ($k = 30, 60, 90$). We can clearly notice that when the wealth and profit information is taken into consideration, the generated clusters for all k values have higher qualities, especially in terms of wealth and profit. This result is expected and confirms that the addition of wealth and profit information can make the wealth features and profit features more uniform and homogeneous. Unexpectedly, the clusters that consider the wealth and profit information also show more consistent decisions. This indicates that there are indeed correlations between decisions and outcomes, and these implicit correlations can be captured by the Doc2Vec model. So in addition to the wealth and profit, the decision consistency is also enhanced. In conclusion, the comparative results reveal that joint consideration of both decision information and wealth-related information will bring better-quality clustering results than the one with only decision information considered.

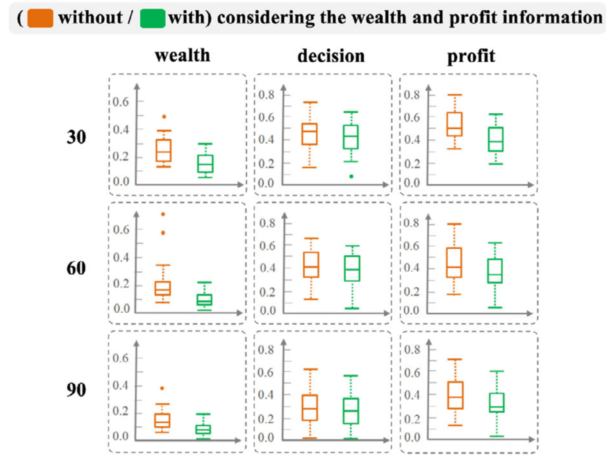

Figure 5: *Boxplots summarize the comparative results for the clusters generated in the Doc2Vec-based embedding space, where a lower score indicates better clustering results.*

5. Economic Behaviour Pattern Visualization

To enable users to easily capture the economic behaviour patterns and their correlations to the outcomes, we design a visualization system in which a control panel (Figure 1a) and four coordinated visualization views are provided including Decision View (Figure 1b), Outcome View (Figure 1c), Comparison View (Figure 1d) and Individual View (Figure 1e).

5.1. Behaviour pattern embedding

Based on the behaviour pattern extraction model in Section 4, in order for the distance between each pair of behaviours to represent their similarity, we employ t-SNE [GCL*22] to project all behaviours into a two-dimensional space, where each point represents a behaviour created by a subject in some experimental round. Thus, similar behaviours (we define them as a pattern **(R1)** in this paper) would be closely distributed, which can be easily detected through visual perception and common clustering methods such as DBSCAN [EKSX96] and K-Means [HW79].

In the Decision View, a couple of colour encoding schemes are designed to reveal the correlation between behaviour patterns and the initial wealth and profits **(R3)**. For the initial wealth encoding, the green colour means that the behaviour is created by a rich subject before he enters the round, while the red colour means a poor subject' s behaviour. For the profit encoding, the green colour denotes a positive profit while the red colour means a negative profit. Points encoded with wealth or profit would help the users easily identify those clusters and outliers of interest. Thus, users can easily specify a group of behaviours in the Decision View with the rectangular

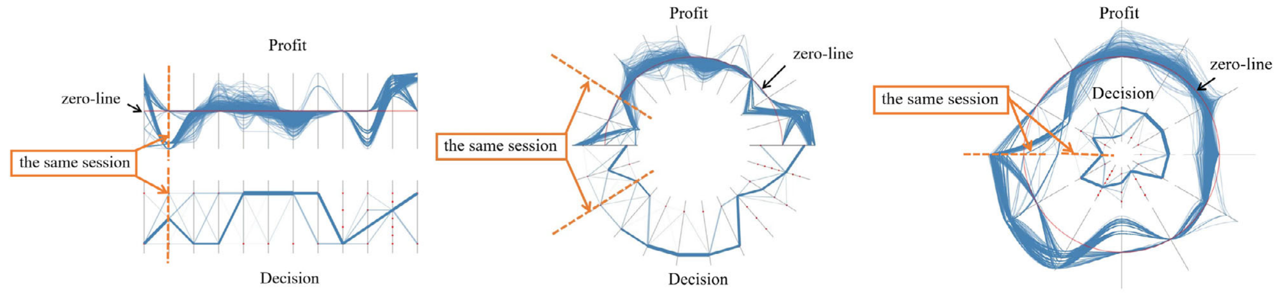

(a) Juxtaposed parallel coordinates (b) Divided radar map (c) Nested radar map

Figure 6: *A couple of design alternatives for displaying patterns in the Comparison View. (a) Arranging two parallel coordinates vertically to respectively show the profit information and decision information. (b) Dividing a radar map into an upper part and a lower part to respectively show the profit information and decision information. Our system uses a nested radar map (c) to combine two kinds of information tightly, providing a better overview of the behaviour patterns.*

marquee tool, to define a behaviour pattern of interest. Besides, the control panel (Figure 1a) offers a timeline and checkboxes to facilitate user interaction. The period can be specified by brushing the timeline, and multiple sessions of interest can be selected by ticking the corresponding check boxes, which can trigger interaction in the Decision View with the desired behaviours re-learned by the behaviour pattern extraction model.

5.2. Pattern visualization and comparison

The Comparison View provides the overviews of multiple behaviour patterns, which are those specified in the Decision View. As shown in Figure 1(d), the side-by-side comparisons of behaviour patterns are provided by dividing the Comparison View into horizontally juxtaposed regions from left to right [YBL*19]. Each region is dedicated to a pattern, thus users can quickly explore the difference in terms of decision preference and economic profit between patterns **(R2)**.

Figure 1(d) provides an overview of a behaviour pattern, which is composed of a line chart in the upper part and a nested radar map in the lower part. The line chart has a horizontal axis for rounds and two vertical axes for the number of subjects and mean profit. In the background, the height of the theme river, corresponding to the left vertical axis, encodes the number of subjects whose behaviours belong to this pattern. In the foreground, the solid line encodes the mean profit resulting from the behaviours of this pattern and the horizontal dotted line represents the zero-profit baseline, which all correspond to the right vertical axis. The nested radar map consists of an inner and outer part, which is constructed with 12 co-centrical circular axes. Each axis represents one experimental session. The outer part displays the dynamic profits changed with sessions, where the blue line representing a behaviour is drawn by connecting the profit values on 12 axes. A red circle crossing all axes is taken as the zero-profit baseline. The point outside of the baseline encodes the positive profit. The larger the profit value, the longer distance it is from the baseline. In a symmetrical manner, the point inside of, and distant towards the baseline encodes a large negative profit. In the inner part, discrete nodes of each axis represent the available choices in the corresponding session. The width and colour depth of the line between adjacent sessions encode the flow between two decisions in adjacent sessions. The names of the dominant decisions in each session are displayed in the middle position between the inner part and the outer part.

In Figure 6, a couple of possible designs would arrange two parallel coordinates vertically or divide a radar map into an upper part and a lower part, to show the profit information and decision information. However, the continuous profit values and discrete decision options create a significant difference, making it difficult for users to perceive the graph as a cohesive whole. Our adopted design applies a nested radar map to combine the two features more tightly, offering an intuitive overview of behaviour patterns and clearly revealing the impact of decisions on session-level profit.

5.3. Correlation between pattern and wealth

The Outcome View shows the wealth variables for different behaviour patterns, enabling users to explore the correlation between behaviour patterns and wealth intuitively **(R3)**. The experts pay more attention to these two variables, namely the initial wealth and the profits, which are both key factors in exploring the mutual effect between behaviour patterns and wealth. This is because the former can help explain how wealth affects behaviours, and the latter can illustrate how the behaviours affect wealth. We design the Outcome View as a Cartesian coordinates. As shown in Figure 1(c), x-axis represents the initial wealth values while y-axis represents the profit values. The users can change the axes of x-axis and y-axis according to their requirements. Each point presents a behaviour, which is placed in the Cartesian coordinates according to the corresponding initial wealth and profit. By observing point distribution, users can quickly see how behaviour patterns and wealth affect each other. In addition, we provide a clustering mode (Figure 7a) to establish the correlation between the behaviour pattern and wealth from a high level. In this case, all points of the same pattern are merged into a circle, which is placed at the centroid of these points. Meantime, the radius of this circle encodes the number of points. The circle represents statistical information about the behaviour pattern, including average initial wealth, average profit and total number

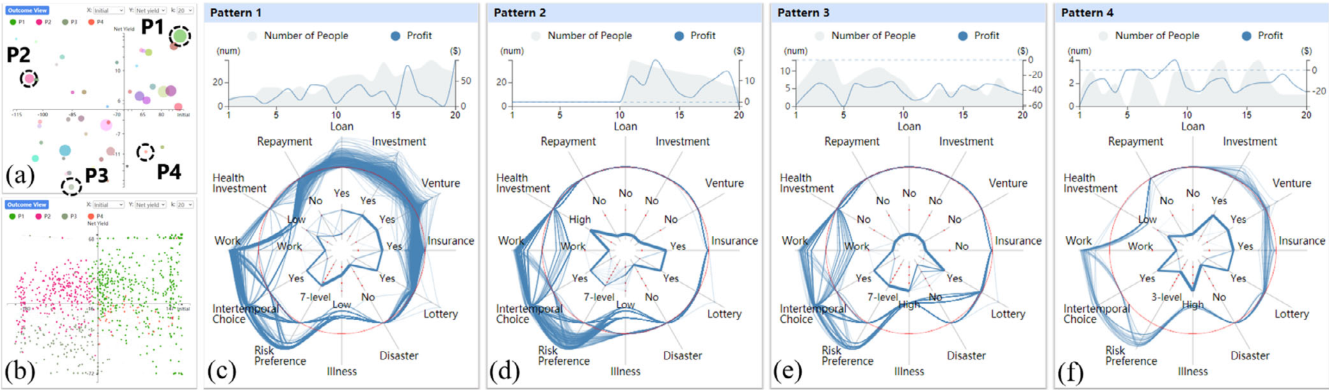

Figure 7: *Four behaviour patterns are displayed in the Outcome View under the mode of (a) cluster and (b) scatter, which is denoted as P1 (the rich subjects earned wealth), P2 (the poor subjects earned wealth), P3 (the poor subjects lost wealth) and P4 (the rich subjects lost wealth). The Comparison View shows the pattern overviews in terms of decision combinations and dynamic profits in (c)–(f).*

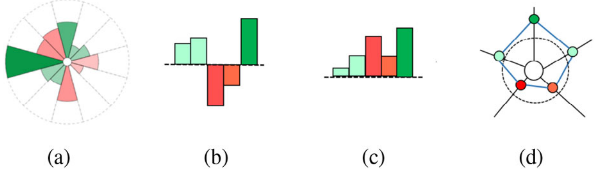

Figure 8: *Four design alternatives for displaying the profits of a subject in the same round. Red colour means a negative profit, and green colour means a positive profit. (a) is utilized to display profits in our system. (b) and (c) are linear designs, which need more space to achieve a similar level of legibility. (d) is a star glyph, in which the lines are difficult to visually capture especially when the glyph is small and has a low colour saturation.*

of behaviours. Thus, users can quickly identify patterns of interest, *e.g*. which behaviour pattern can help subjects gain the most profits.

5.4. Individual behaviour pattern

The Individual View displays the detailed information of individuals at the micro-level, which helps users to further track and compare different subjects' decision strategies and wealth accumulations over time **(R4)**. As shown in Figure 1(e2), the horizontal and vertical axes, respectively, encode the experimental rounds and the total wealth. Initially, 20 whisker plots are displayed to show the distribution of wealth values of all subjects in each round. If one subject is selected through other views, his/her wealth-change line will be shown. To distinguish each session's contribution to the net profit of one round, a petal glyph [KAW*14] is designed to overlay the corresponding coordinate point. The glyph consists of 12 sectors representing the 12 sessions (Figure 8a), of which the angles are consistent with the nested radar maps. Red and green indicate a negative profit and a positive profit, and the length of each sector encodes the absolute value of the profit. In the beginning of glyph design, we considered several possible designs (Figures 8b–d) such as bar chart and star glyph. However, the bar chart requires more space to display, thus achieving a comparable level of legibility as that of the petal glyph [MR10]. The single line in the star glyph is relatively difficult to perceive especially when the glyph is small and has a low colour saturation. Ultimately, we opted to discard these options in favour of the current design, which is more compatible with other views and easy to encode and convey numerical profit values [ZWC*17].

A list of horizon graphs is shown on the left side of the Individual View (Figure 1(e1)). Each horizon graph represents a subject, which summarizes the time-varying profits of all sessions. Similarly, red and green colours indicate negative and positive profits, and their intensity indicates the magnitude. The horizon graph and the glyph-embedded lines convey the information of one subject from different angles. Beyond showcasing dynamic wealth, glyph-embedded lines facilitate comparing the profits of different sessions in the same round. In contrast, the horizon graph provides a more intuitive comparison on the profits of one session in different rounds.

6. Evaluation

Three experts E1, E2 and E3 are invited to employ *EBPVis* for the evaluation of the capabilities of our system. E1 is an experienced researcher specializing in experimental economics, contributing 6 years of dedicated research at the Center for Economic Behavior and Decision-making as mentioned in Section 3. E2, a microeconomics professor, has substantial expertise in analysing income distribution among diverse individuals and households. E3 fulfils the role of a dedicated government poverty alleviation worker, committed to operating in impoverished areas and overseeing the meticulous implementation of targeted policies.

6.1. Correlation analysis between patterns and outcomes

E1 explored the behaviour data covering all the sessions and all rounds, hoping to find some correlations between patterns and outcomes **(R3)**. First, he navigated the patterns under the cluster mode in the Outcome View, as shown in Figure 7(a). Then, four patterns (denoted as P1, P2, P3 and P4) in the first, second, third and fourth

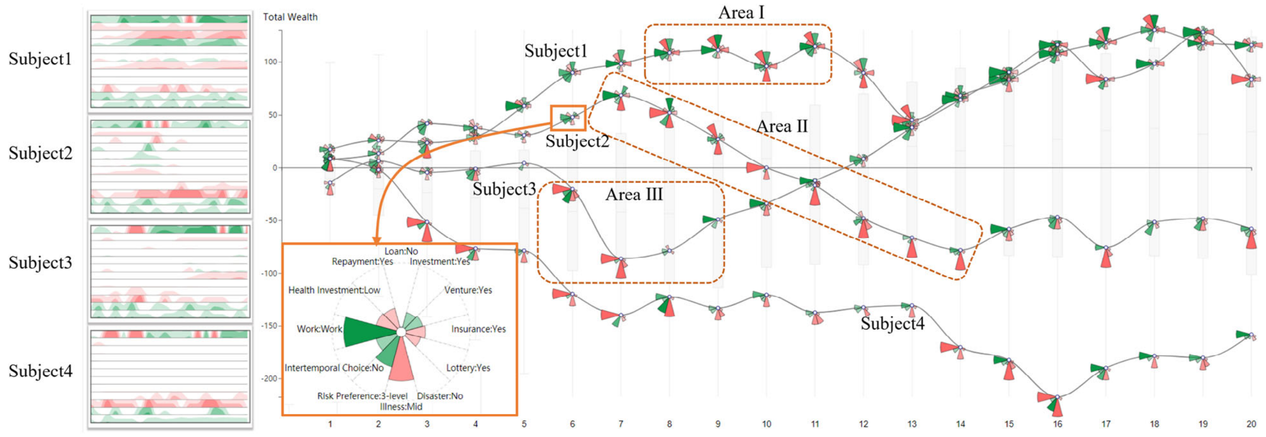

Figure 9: *Four interesting subjects are highlighted in the Individual View. They have different wealth evolutionary patterns due to their own effort, luck and decisions. The four horizon graphs are listed on the left to represent the patterns of their whole life.*

quadrants were selected, and their pattern overviews were displayed in the Comparison View at the same time. E1 said that the selected patterns corresponded to four different behaviours: the rich subjects earned wealth, the poor subjects earned wealth, the poor subjects lost wealth and the rich subjects lost wealth. To further understand the reasons behind these patterns, he examined the Comparison View to explore the differences between these patterns in depth **(R2)**. For the pattern P1 (Figure 7c), a large number of lines were located outside of the red circle along the axes of loan, investment and venture, indicating that the three sessions all brought in respectable profits for the subjects.

Besides, by observing the line chart at the top, E2 found the number of subjects belonging to this pattern gradually increased. E2 thought that repeated experiments gave subjects the opportunity to learn and adapt their behaviour to experience, and such decision behaviour of 'earning money' was learned by more and more people. Next, E2 paid attention to the pattern P2, which was beyond our understanding. Figure 7(d) showed a 'High' label on the axis of 'Health investment' and a 'Yes' label on the axis of 'Insurance', indicating that the subjects purchased health insurance and property insurance. However, there was no reduction in their wealth because the corresponding lines overlapped with the red circle along the two axes. Our experts told us that to help the poor get rid of poverty, a poverty alleviation policy was introduced after the tenth round to provide cost-free health insurance and property insurance for a part of poor subjects. They speculated that these subjects were the lucky ones enjoying the poverty alleviation policies.

E3 examined the line chart and found this pattern started to appear in the 11th round. He said such kind of pattern was successfully discovered by our system. Furthermore, E3 believed that the poverty alleviation policy prevented the poor subjects from property loss caused by illness and disaster, resulting in their wealth accumulation. In contrast, poor subjects in pattern P3 (Figure 7e) lacked health and property insurance, leading to significant wealth reduction from illness and disaster losses outweighing earnings. Finally, the pattern P4 was explored by the experts. As shown in Figure 7(f), though the subjects made investments and venture investments, the profits were relatively low compared with pattern P1. What's worse, the subjects had a 'High' illness, which meant they had to spend more money on getting better treatment. All these factors had a negative impact on their wealth. E2 commented that such behaviours were uncommon, as few patterns appeared in the fourth quadrant of the Outcome View.

6.2. Individual behaviour exploration

After exploring a number of subjects, E1 found four interesting individual cases **(R4)** as shown in Figure 9. It was shown that Subject 1 with a high working ability constantly borrowed money to invest and re-pay the loan as shown in Area I of Figure 9. But he had been unlucky in the session of the venture, his wealth fluctuated wildly. In contrast, the wealth of Subject 2 continued to fall due to the illness from the 7th round to the 14th round, and dropped from the middle class to the poor as shown in Area II of Figure 9. For Subject 3, he lost the job in the 6th round and the 7th round as shown in Area III of Figure 9. What's worse, he was stricken by a serious illness in the 7th round. But later this subject had been working hard and made him rich again. As we can see, the leftward sectors, corresponding to the work session, dominate the glyphs obviously from the 9th round to the end. Compared with four subjects, Subject 4 was a poor person throughout the entire experiment. Though he struggled to increase his wealth by working in several rounds, he still failed to escape poverty. Besides, the horizon graphs on the left side of Figure 9 also revealed the corresponding wealth evaluations clearly, where green and red represented the poor and rich. E1 commented 'This supports the theory of the vicious cycle of poverty, asserting that the poor are getting poorer and poorer, and the rich are getting richer and richer. However, if a poor person can seize the opportunity and make a change actively, he will achieve success'.

6.3. Interactive pattern extraction by disrupting the session order

Events occurring in real-world scenarios rarely adhere to a consistent order for individuals. To authentically simulate real-world situations, we meticulously developed a web-based system that allows

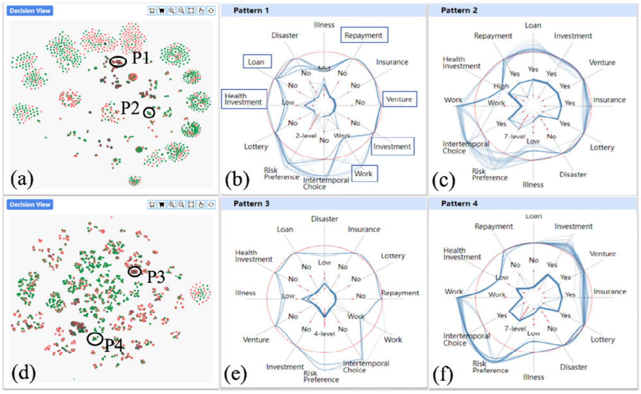

Figure 10: *The exploration of the session order on behaviour patterns. (a) The decisions in the first six sessions are taken into consideration for the projection. (b) and (c) are the behaviour patterns of P1 (randomized) and P2 (default) in (a), respectively. (d) All session decisions are taken into consideration. (e) and (f) show behaviour patterns of P3 (randomized) and P4 (default) in (d).*

for a randomized session order, faithfully mirroring the original system employed by our collaborative team for experiments. We engaged 30 undergraduate students (17 males and 13 females) in 20 rounds of experiments, each with randomized session orders. Subsequently, their experimental results were integrated with the initial experimental data into the *EBPVis*.

E1 began by selecting the decisions on the first six sessions (Work, Health Investment, Repay, Loan, Investment, and Venture) through the Control Panel and clicking the 'DBSCAN' button to observe the embedding result in the Decision View, as shown in Figure 10(a). Behaviours in a randomized order are marked with black strokes to distinguish them from the default order. E1 picked two clusters: P1 with a randomized order, and P2 with a default order. Figures 10(b) and 10(c) provided pattern overviews within these clusters. He observed a consistent line distribution in the radar maps of the two figures in the first six sessions, whereas the remaining sessions displayed more scattered lines. This indicates that Doc2Vec can effectively capture similar behavioural patterns. In addition, he noticed that most of the projection points in P1 are filled in red, while most of the points in P2 are filled in green. This suggests that randomized orders are associated with wealth loss, while default orders tend to result in wealth gain. To explain this phenomenon, upon careful pattern analysis, E1 observed that compared to the experiments in P2 which consciously placed the process of accumulating and safeguarding wealth in the first half of fixed orders, sessions in P1 featured disasters and illness early, with work coming later. Such orders potentially made individuals with limited funds for insurance, lottery tickets and investments. It was challenging for them to mitigate losses from disasters and illnesses, consequently missing out on opportunities to become rich.

E2 proceeded by concurrently selecting all decisions via the Control Panel, promptly updating the Decision View. As depicted in Figure 10(d), the projection points exhibited a distribution pattern featuring sizable clusters that were further segmented into smaller ones. He deduced that this resulted from fewer similar patterns aligning across 12 decision-making rounds compared to the case where only six rounds were required. Among them, he chose cluster P3 with randomized order and P4 with default order, examining their pattern details in Figures 10(e) and 10(f). Notably, regardless of order disruption, he observed consistent behaviours within the radar map at all sessions, with outer also maintaining consistency. Through the above comparison and explorations, experts believed that our pattern extraction model could help them identify behavioural patterns with common characteristics of both decisions and profits on the sessions of their interests **(R1)**, and verify the impact of experimental order on decision-making.

6.4. Expert interview

System capability and effectiveness. The three experts had never used such a visual analytics tool before, and they were all impressed by our system. E1 told us, 'Although traditional statistics analytical methods are useful, this interactive system provides more flexibility and feedback to explore, display and understand the economic behaviors'. In addition, all experts thought the Decision View was helpful for them to identify and extract the economic behaviour patterns. E1 noted that while traditional tools like Stata and SPSS rely on pre-defined models and manual feature extraction, our Doc2Vec representation captures subtle behaviour similarities and differences often missed by these traditional methods. Particularly, E3 told us that the multiple attribute combination and the timeline were useful to specify attributes and periods of interest and extract the relevant patterns. E2 stated that the success of economic theories relies on explaining market participant behaviour. He praised our system for effectively elucidating subjects' decision–profit interactions, revealing behavioural heterogeneity and facilitating the quick comparison of different behaviour patterns. However, E2 also mentioned that since he took part in the economic experiment, he wanted to know and explore his decision behaviours. So he suggested that if the system could allow users to retrieve one subject and further highlight his decision data in different views, it would be more user-friendly. More importantly, he hopes that future versions of *EBPVis* can provide exploratory functions guided by different experimental economic tasks, such as analysing the impact of immediate effects, endowment effects and learning effects on individual preferences.

Visual design and interactions. The visualization views were appreciated by our experts. They found the Outcome View designed in the system was visually appealing and conveyed complementary information regarding the outcome variables. E2 commented that the functions of each view and their interactions enabled him to explore the behaviour patterns at different levels. As for the nested radar map, E1 commented, 'Its dual-layered design within a limited visual space, allowing for detailed analysis of complex behavior patterns, making it easier for me to compare decision and profit information simultaneously'. E3 found the Individual View and the glyph design exciting and easy to understand, 'I can filter a few subjects of interest through this view and take them as a good or bad example to analyze whether these subjects actually learn from their experiences'. Moreover, he felt it a little tricky to track one subject due to the crossings and overlaps between curves and glyphs when many subjects were displayed in the Individual View. So, he thought

Table 1: *User study results.*

Category	User task	Success rate	Avg time (s)	Std dev time (s)
Behaviour pattern identification	1.1 Which pattern can help subjects gain the most profits?	100%	14.7	2.9
	1.2 Which pattern can make subjects lose the most profits?	100%	12.0	2.1
	1.3 Which pattern can make the poor subjects gain profits?	100%	10.1	2.2
	1.4 Which pattern can make the rich subjects lose profits?	100%	9.8	1.9
Behaviour patterns comparison	2.1 Which session has the most consistent decisions for a pattern?	95%	22.4	5.3
	2.2 Which session has the most profits for a pattern?	100%	25.6	5.7
	2.3 Which session has the largest difference in profits between two patterns?	100%	39.4	8.2
	2.4 Which session has the same decision but opposite profits between two patterns?	95%	44.9	9.1
Individual behaviour exploration	3.1 Which round has one subject been losing the most profits?	95%	10.1	2.6
	3.2 Which session has the biggest impact on one round for a subject?	85%	17.2	4.5
	3.3 Which session has one subject been making money in?	90%	18.3	5.1
	3.4 Which session has one subject been losing money in?	90%	16.7	4.8

it would be better to highlight one specific subject in the context of multiple subjects. E1 suggested a preference for a simplified system to improve user experience, indicating that the current system is overly complex. He added, 'For a novice, it might not be easy to use it right away. Simplifying the interface by filtering the most important elements is recommended'.

6.5. User study and feedback

To further evaluate the usability of our system, we invited 20 undergraduate and graduate students (12 males and eight females) majoring in economics and computer science to participate in the user study. They were all trained to use our system until they were familiar with the system. Typically, users only need 15–30 minu of training time to understand the meaning of each view and the function of our system. Thereafter, they were asked to answer questions by using the system within the specified time. We designed a series of questions, which are closely related to the analysis tasks in Section 3. For each question, we set the maximum completion time as 60 s. The percentage of users who were able to complete tasks correctly within the given time, the average and standard deviation of completion times were recorded, as shown in Table 1.

We collected their feedback on the system when they were performing the above tasks. They all thought the system was very novel and interesting, which could help them more intuitively explore economic behaviours, whether from the overall or micro perspective. Among them, the undergraduate students in economics told us that they had taken experimental economics-related courses before, but the involved principles seemed to be very theoretical and abstract. 'The behavioral patterns found in this system can work as a good example to explain what subjects do in the experimental process. User preferences, behavioral differences and decision-making rules are revealed intuitively. This system can be used as an auxiliary teaching tool to enhance students' comprehension of the concept of experimental economics', they commented. Besides, some students, who had participated in the experiment, said that when they faced different tasks in the virtual experimental environment, they often decided from their own experience. They mentioned: 'Some results provided by your system are useful and inspiring. For example, the non-rational behaviors derived from the comparison be-

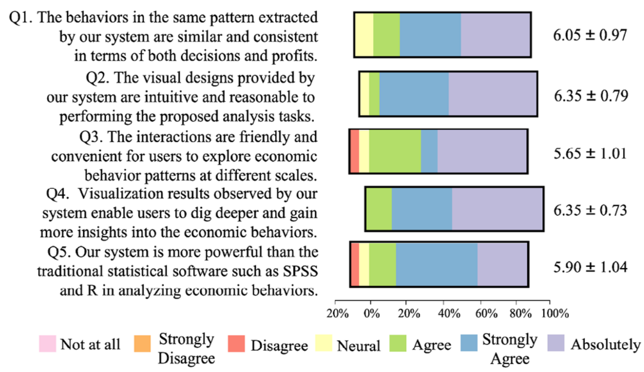

Figure 11: *Description and results of questionnaires. The right column denotes mean ± SD.*

tween different behavior patterns let us rethink our decisions and strategies. We are sure that we will receive more rewards if given another chance to participate in the experiment'. Despite the positive feedback, our participants identified some limitations of *EBPVis* and offered suggestions for improvement. One user out of 20 did not complete UT2.1 and UT2.4 on time, as the user thought that the inner radar map for visualizing decisions was relatively small, so he had to spend more time observing the details. We observed a decline in success rate and large variations starting from UT3.2. An outlier in UT3.2 noted that identifying which session each petal corresponded to is a little distracting, indicating that resolving UT3 would require more user interaction and observation among visual components. He also pointed out that displaying multiple subjects simultaneously in the Individual View might cause petal glyphs to overlap. Another participant, who did not complete UT3.3 and UT3.4, mentioned unfamiliarity with horizon graphs and difficulty in determining which session each layer represented. Hence, he suggested adding text labels to the horizon graphs to clarify the corresponding experimental sessions.

To further gather more explicit feedback on our system, we put forward five questions covering behaviour patterns extraction, visual design, interaction, case study and other software comparisons, using a 7-point Likert scale. As shown in Figure 11, the answers to the five questions are generally positive and affirmative. Among them, our behaviour pattern extraction (Q1), visual design (Q2) and case study (Q4) all obtained an average score of more than 6 points out of 7, which was very encouraging. However, in Q2, some participants without a visualization background reported needing more effort to understand certain visual designs, such as nested radar maps and horizon graphs. Many participants initially struggled to remember all the visual encodings and interactions, which required additional explanations. Yet, after a few trials, they were able to quickly and successfully gain the desired insights. The interaction (Q3) garnered more mixed reviews. Some users appreciated the range of functions, while others felt it lacked innovative features. This feedback highlights a potential area for improvement in integrating more intuitive and novel interactions. Our system (Q5) was broadly appreciated, although we received one criticism that our system has more advanced visual designs than traditional statistical systems, but it usually needs some prior knowledge to get started. Although our system offers significant advantages in visual analysis and behaviour pattern recognition, simplifying the initial learning curve could make it more widely accessible to a broader user base.

7. Discussion

Experiment designers in the Center for Economic Behavior and Decision-making appreciated the insights and findings provided by our system. However, our system does have several limitations.

Scalability. Analysing system scalability ensures researchers can apply *EBPVis* to both pilot and extensive studies. Doc2Vec converts large behaviour sequences into continuous vectors without increasing linear computational complexity and supports incremental learning, allowing the model to grow with the data. From a visualization perspective, users in the case study typically compare 2–4 patterns, each with 10–30 records. The comparison view effectively meets these needs by displaying up to five patterns side-by-side, accommodating dozens of records. However, viewing more patterns simultaneously requires scrolling, which can increase cognitive load. Sampling and filtering can alleviate visual clutter in scatter plots and nested radar maps for larger datasets, supporting more complex queries. When examining many subjects in the Individual View, the visual clutter resulting from the overlaps of glyphs might arise. To address this, we plan to integrate methods such as difference views that directly encode differences among multiple subjects and patterns, posing new challenges for screen space usage and layout. Furthermore, our system records decisions and earnings for each subject per session. When the number of sessions exceeds 20, it will lead to small sector angles in the nested radar and petal glyph, which will hinder perception. To mitigate this, we can allow users to select specific sessions of interest for further exploration.

Cognitive load of visual design. Our system has received widespread acclaim from experts and users alike, but there are some reservations. A key point is that while our advanced visual design surpasses traditional statistical methods, it requires a certain level of prior knowledge to use effectively. This highlights a critical challenge: balancing the intuitiveness and expressiveness of the visualization. To ease the initial learning curve and address this issue, we plan to add more annotations and interactions in future work. For instance, providing dynamic tooltips with detailed values and explanations when users hover over or click on specific sessions of the petal glyph, and highlighting corresponding sessions of different petals to help users quickly distinguish and compare them.

Behaviour pattern representation. Experts unanimously agree that Doc2Vec is highly effective in representing behaviour patterns compared to traditional descriptive and inferential statistical methods. Its exceptional ability to capture contextual information aids in distinguishing the combination of decision and profit, making it suitable for highly interpretative economic research. However, Doc2Vec is less effective than pre-trained models in predicting subjects' behaviours and wealth changes, as it struggles to capture complex non-linear relationships. To address this, we propose integrating pre-trained models [WLW*23] such as BERT [DCLT19] and GPT-4 [RNS*18]. BERT's bidirectional training and GPT-4's extensive pre-training enable these models to capture deeper contextual nuances. By using Doc2Vec for initial feature extraction and then fine-tuning BERT and GPT on our dataset, we can create a hybrid model. This approach leverages Doc2Vec's contextual embeddings and the predictive power of BERT and GPT, enhancing the accuracy of unknown behaviour and wealth predictions. Integrating these functionalities into an all-in-one platform would significantly enhance efficiency for experimental economics researchers.

Generalizability. As Vernon Smith, the father of experimental economics, stated, this field typically employs repeated rounds of experiments to study the interactive decision-making behaviour of subjects driven by specific factors [Smi06]. *EBPVis*, though designed for specific experimental data, can be easily adapted for other experimental economics applications with minor modifications. For instance, in transportation, one could analyse the impact of travel costs on transportation mode choices by simulating various scenarios (*e.g*. peak hours, traffic congestion) and collecting travel choices and cost data. In medical decision-making, one could evaluate doctors' altruistic behaviour by designing various scenarios (*e.g*. resource-limited emergencies, patient financial situations) and collecting their decision data. These data can be utilized and transmitted through our system with modified visual components, such as glyphs. Additionally, our system's potential is not limited to experimental economics. It can be extended to address more general issues of behaviour sequence and outcome analysis in repeated sessions and rounds, such as analysing students' behaviour patterns and learning outcomes. Our future goal is to integrate an attribute editor for flexible data attribute definition and modification.

8. Conclusion

In this work, we develop a visual analytics system, *EBPVis*, enabling economists to visually explore human behaviour patterns. We utilize a Doc2Vec model to transform the economic behaviours into an embedding space, in which the behaviour patterns can be easily perceived and extracted by users. To further explore the correlation between decisions and outcomes, a set of visual designs such as the Outcome View, the Comparison View and the Individual View are designed for users to visually explore the human behaviour patterns from multiple perspectives. Case studies, expert feedback and user studies based on a real-world dataset have demonstrated the effectiveness and practicability of our system in the representation of economic behaviour patterns and certification of economic theories.

In future work, we will collaborate with experts to conduct long-term research, aiming to further evaluate the usability and effectiveness of *EBPVis* in analysing behaviour patterns. Additionally, we plan to integrate advanced AI techniques to enhance the capture and prediction of complex behaviour patterns. Our ultimate goal is to extend the system's application to a broader range of experimental economics scenarios, supporting diverse research needs.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (Nos. 62277013 and 62177040), the National Statistical Science Research Project (No. 2023LZ035), the Key R&D 'Jianbin' Tackling Plan Programme in Zhejiang Province, China (No. 2023C01119), the Public Welfare Plan Research Project of Zhejiang Provincial Science and Technology Department (Nos. LTGG23H260003 and LGF22F020034), Open Project Programme of the State Key Laboratory of CAD&CG (No. A2301) and Zhejiang Provincial Natural Science Foundation of China (No. LTGG24F020006).

Conflicts of Interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

References

[ABRS21] AUCLERT A., BARDÓCZY B., ROGNLIE M., STRAUB L.: Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica 89* (Jan. 2021), 2375–2408. doi: 10.3982/ECTA17434.

[AS*94] AGRAWAL R., SRIKANT R.: Fast algorithms for mining association rules. In *Proceedings of 20th International Conference on Very Large Data Bases, VLDB* (1994), vol. *1215*, Citeseer, pp. 487–499.

[BKC*13] BORGO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics (State of the Art Reports)* (2013), pp. 39–63. doi: 10.2312/conf/EG2013/stars/039-063.

[Bol10] BOLTON G.: Testing models and internalizing context: A comment on "theory and experiment: What are the questions?". *Journal of Economic Behavior & Organization 73*, 1 (2010), 16–20. doi: 10.1016/j.jebo.2009.11.002.

[BQL*20] BO Y., QI M., LIU S., CUI J., HAN Y.: Profit distribution and managers' behavior in vertical integrated medical delivery systems—an experimental economics study. *BMC Health Services Research 20*, 1 (2020), 1–9. doi: 10.21203/rs.2.17196/v1.

[BS07] BRÜGGEN A., STROBEL M.: Real effort versus chosen effort in experiments. *Economics Letters 96*, 2 (2007), 232–236. doi: 10.1016/j.econlet.2007.01.008.

[CD17] CARDELLA E., DEPEW B.: Output restriction and the ratchet effect: Evidence from a real-effort work task. *Games and Economic Behavior 107* (Nov. 2017), 182–202. doi: 10.1016/j.geb.2017.11.005.

[CHLS03] CSALLNER C., HANDTE M., LEHMANN O., STASKO J.: FundExplorer: Supporting the diversification of mutual fund portfolios using context treemaps. In *Proceedings - IEEE Symposium on Information Visualization, INFO VIS* (Jan. 2003), pp. 203–208. doi: 10.1109/INFVIS.2003.1249027.

[CT*10] CAMERON A. C., TRIVEDI P. K., et al.: *Microeconometrics Using Stata* (vol. 2). Stata Press College Station, TX, 2010.

[CvW17] CAPPERS B. C., VAN WIJK J. J.: Exploring multivariate event sequences using rules, aggregations, and selections. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 532–541. doi: 10.1109/TVCG.2017.2745278.

[CXR17] CHEN Y., XU P., REN L.: Sequence synopsis: Optimize visual summary of temporal event data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 45–55. doi: 10.1109/TVCG.2017.2745083.

[DC08] DERCON S., CHRISTIAENSEN L.: Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia. *Journal of Development Economics 96* (Jan. 2008), 159–173. doi: 10.1016/j.jdeveco.2010.08.003.

[DP08] DESSAU R. B., PIPPER C. B.: "R"–project for statistical computing. *Ugeskrift for laeger*, *170*, 5 (2008), 328–330.

[DCLT19] DEVLIN J., CHANG M.-W., LEE K., TOUTANOVA K.: BERT: pre-training of deep bidirectional transformers for language understanding. *In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)* (Jun. 2019), pp. 4171-4186. doi: 10.18653/V1/N19-1423.

[DE20] DOVAL L., ELY J. C.: Sequential information design. *Econometrica 88*, 6 (2020), 2575–2608. doi: 10.3982/ECTA17260.

[EKSX96] ESTER M., KRÖGER P., SANDER J., XU X.: A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD* (Jan. 1996), vol. *96*, pp. 226–231.

[FVLK*17] FOURNIER-VIGER P., LIN J. C.-W., KIRAN R. U., KOH Y. S., THOMAS R.: A survey of sequential pattern mining. *Data Science and Pattern Recognition 1*, 1 (2017), 54–77.

[GCL*22] GOVE R., CADALZO L., LEIBY N., SINGER J., ZAITZEFF A.: New guidance for using t-SNE: Alternative defaults, hyperparameter selection automation, and comparative evaluation. *Visual Informatics 6* (Apr. 2022), 87–97. doi: 10.1016/j.visinf.2022.04.003.

[GGJ*21] GUO Y., GUO S., JIN Z., KAUL S., GOTZ D., CAO N.: A survey on visual analysis of event sequence data. *IEEE*

Transactions on Visualization and Computer Graphics (2021). doi: 10.1109/TVCG.2021.3100413.

[GJG*18] Guo S., Jin Z., Gotz D., Du F., Zha H., Cao N.: Visual progression analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2018), 417–426. doi: 10.1109/TVCG.2018.2864885.

[GN15] Gunaratne J., Nov O.: Informing and improving retirement saving performance using behavioral economics theory-driven user interfaces. In *Acm Conference on Human Factors in Computing Systems* (Apr. 2015), pp. 917–920. doi: 10.1145/2702123.2702408.

[GXZ*17] Guo S., Xu K., Zhao R., Gotz D., Zha H., Cao N.: EventThread: Visual summarization and stage analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics 24*, 1 (2017), 56–65. doi: 10.1109/TVCG.2017.2745320.

[GZW*19] Gotz D., Zhang J., Wang W., Shrestha J., Borland D.: Visual analysis of high-dimensional event sequence data via dynamic hierarchical aggregation. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 440–450. doi: 10.1109/TVCG.2019.2934661.

[HGM*97] Hoffman P., Grinstein G., Marx K., Grosse I., Stanley E.: DNA visual and analytic data mining. In *Proceedings. Visualization'97 (Cat. No. 97CB36155)* (1997), IEEE, pp. 437–441. doi: 10.1109/VISUAL.1997.663916.

[HPY00] Han J., Pei J., Yin Y.: Mining frequent patterns without candidate generation. *ACM Sigmod Record 29*, 2 (2000), 1–12. doi: 10.1145/342009.335372.

[HSvK*19a] Hu R., Sha T., van Kaick O., Deussen O., Huang H.: Data sampling in multi-view and multi-class scatterplots via set cover optimization. *IEEE Transactions on Visualization and Computer Graphics* (Aug. 2019), 1. doi: 10.1109/TVCG.2019.2934799.

[HSVK*19b] Hu R., Sha T., Van Kaick O., Deussen O., Huang H.: Data sampling in multi-view and multi-class scatterplots via set cover optimization. *IEEE Transactions on Visualization and Computer Graphics 26*, 1 (2019), 739–748. doi: 10.1109/TVCG.2019.2934799.

[HW79] Hartigan J. A., Wong M. A.: A k-means clustering algorithm. *Applied Statistics 28*, 1 (1979).

[HW13] Heinrich J., Weiskopf D.: State of the art of parallel coordinates. In *Eurographics (State of the Art Reports)* (2013), pp. 95–116.

[Kan00] Kandogan E.: Star coordinates: A multi-dimensional visualization technique with uniform treatment of dimensions. In *Proceedings of the IEEE Information Visualization Symposium* (2000), vol. *650*, Citeseer, pp. 22.

[KAW*14] Ko S., Afzal S., Walton S., Yang Y., Chae J., Malik A., Jang Y., Chen M., Ebert D.: Analyzing high-dimensional multivariate network links with integrated anomaly detection, highlighting and exploration. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), IEEE, pp. 83–92. doi: 10.1109/VAST.2014.7042484.

[KCA*16] Ko S., Cho I., Afzal S., Yau C., Chae J., Malik A., Beck K., Jang Y., Ribarsky W., Ebert D.: A survey on visual analysis approaches for financial data. *Computer Graphics Forum 35* (June 2016), 599–617. doi: 10.1111/cgf.12931.

[KHD*10] Keim D. A., Hao M. C., Dayal U., Janetzko H., Bak P.: Generalized scatter plots. *Information Visualization 9*, 4 (2010), 301–311. doi: 10.1057/ivs.2009.34.

[LLMB18] Law P.-M., Liu Z., Malik S., Basole R. C.: Maqui: Interweaving queries and pattern mining for recursive event sequence exploration. *IEEE Transactions on Visualization and Computer Graphics 25*, 1 (2018), 396–406. doi: 10.31219/osf.io/tg7w5.

[LM14] Le Q., Mikolov T.: Distributed representations of sentences and documents. In *International Conference on Machine Learning* (2014), PMLR, pp. 1188–1196.

[LMW*16] Liu S., Maljovec D., Wang B., Bremer P.-T., Pascucci V.: Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics* (Dec. 2016), 1. doi: 10.1109/TVCG.2016.2640960.

[LVML07] Lenca P., Vaillant B., Meyer P., Lallich S.: Association rule interestingness measures: Experimental and theoretical studies. In *Quality Measures in Data Mining*. Springer, Berlin (2007), pp. 51–76. doi: 10.1007/978-3-540-44918-8_3.

[LWCC17] Liao H., Wu Y., Chen L., Chen W.: Cluster-based visual abstraction for multivariate scatterplots. *IEEE Transactions on Visualization and Computer Graphics 24*, 9 (2017), 2531–2545. doi: 10.1109/TVCG.2017.2754480.

[LWD*16] Liu Z., Wang Y., Dontcheva M., Hoffman M., Walker S., Wilson A.: Patterns and sequences: Interactive exploration of clickstreams to understand common visitor paths. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2016), 321–330. doi: 10.1109/TVCG.2016.2598797.

[MCCD13] Mikolov T., Chen K., Corrado G., Dean J.: Efficient estimation of word representations in vector space. In *Proceedings of Workshop at ICLR 2013* (Jan. 2013).

[MKW24] Munz-Körner T., Weiskopf D.: Exploring visual quality of multidimensional time series projections. *Visual Informatics 8* (May 2024), 27–42. doi: 10.1016/j.visinf.2024.04.004.

[MR10] McGuffin M., Robert J.-M.: Quantifying the space-efficiency of 2D graphical representations of trees. *Information Visualization 9* (June 2010), 115–140. doi: 10.1057/ivs.2009.4.

[MV05] Mosley P., Verschoor A.: Risk attitudes in the vicious circle of poverty. *The European Journal of Development Research 17* (Mar. 2005), 59–88. doi: 10.1080/09578810500066548.

[Org15] ORG Z.: *Handbook of Experimental Economic Methodology*. OUP Catalogue, New York, 2015.

[PWLM22] PAN X., WU L., LONG F., MA A.: Exploiting user behavior learning for personalized trajectory recommendations. *Frontiers of Computer Science 16* (Feb. 2022), 163610. doi: 10.1007/s11704-020-0243-2.

[RLS*18] ROBERTS R., LARAMEE R., SMITH G., BROOKES P., DCRUZE T.: Smart brushing for parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics* (Feb. 2018), 1. doi: 10.1109/TVCG.2018.2808969.

[RNS*18] RADFORD A., NARASIMHAN K., SALIMANS T., SUTSKEVER I.: Improving language understanding by generative pre-training. *OpenAI Blog*, 2018.

[RSE09] RUDOLPH S., SAVIKHIN A., EBERT D.: FinVis: Applied visual analytics for personal financial planning. In *VAST 09—IEEE Symposium on Visual Analytics Science and Technology, Proceedings* (Jan. 2009), pp. 195–202. doi: 10.1109/VAST.2009.5333920.

[SLFE11] SAVIKHIN A., LAM H., FISHER B., EBERT D.: An experimental study of financial portfolio selection with visual analytics for decision support. In *2011 44th Hawaii International Conference on System Sciences (HICSS)* (Feb. 2011), IEEE, pp. 1–10. doi: 10.1109/HICSS.2011.54.

[Smi06] SMITH V.: *Papers in Experimental Economics*. Cambridge University Press, Cambridge, 2006. doi: 10.1017/CBO9780511528354.

[SPS90] SPSS, Inc.: *SPSS/PC Trends: for the IBM PC/XT/AT and PS/2*. SPSS, Chicago, 1990.

[TB99] THALER R., BENARTZI S.: Risk aversion or myopia? Choices in repeated gambles and retirement investments. *Management Science 45* (Mar. 1999), 364–381. doi: 10.1287/mnsc.45.3.364.

[TCS19] TIAN Y., CHIU Y.-C., SUN J.: Understanding behavioral effects of tradable mobility credit scheme: An experimental economics approach. *Transport Policy 81* (2019), 1–11. doi: 10.1016/j.tranpol.2019.05.019.

[TT12] TREMBLAY V. J., TREMBLAY C. H.: *New Perspectives on Industrial Organization. Springer Texts in Business and Economics*. Springer, New York, 2012. doi: 10.1007/978-1-4614-3241-8_19.

[WBK19] WEIMANN J., BROSIG-KOCH J.: *Methods in Experimental Economics. Springer Texts in Business and Economics*. Springer, Cham, 2019.

[WKMD21] WESSLEN R., KARDUNI A., MARKANT D., DOU W.: Effect of uncertainty visualizations on myopic loss aversion and equity premium puzzle in retirement investment decisions. *IEEE Transactions on Visualization and Computer Graphics* (July 2021). doi: 10.1109/TVCG.2021.3114692.

[WLW*23] WEN M., LIN R., WANG H., YANG Y., WEN Y., MAI L., WANG J., ZHANG H., ZHANG W.: Large sequence models for sequential decision-making: A survey. *Frontiers of Computer Science 17*, 6 (2023), 176349. doi: 10.1007/s11704-023-2689-5.

[YBL*19] YUE X., BAI J., LIU Q., TANG Y., PURI A., LI K., QU H.: sPortfolio: Stratified visual analysis of stock portfolios. *IEEE Transactions on Visualization and Computer Graphics* (August 2019) doi: 10.1109/TVCG.2019.2934660.

[ZMZ*21] ZHOU Z., MA Y., ZHANG Y., LIU Y., YANG X., LIU Y.: Context-aware visual abstraction of crowded parallel coordinates. *Neurocomputing 459* (May 2021). doi: 10.1016/j.neucom.2021.05.005.

[ZWC*17] ZHAO X., WU Y., CUI W., DU X., CHEN Y., WANG Y., LEE D., QU H.: SkyLens: Visual analysis of skyline on multi-dimensional data. *IEEE Transactions on Visualization and Computer Graphics* (Aug. 2017). doi: 10.1109/TVCG.2017.2744738.

[ZYC*24] ZHOU Z., YE L., CAI L., WANG L., WANG Y., WANG Y., CHEN W., WANG Y.: ConceptThread: Visualizing threaded concepts in MOOC videos. *IEEE transactions on visualization and computer graphics* (Feb. 2024). doi: 10.1109/TVCG.2024.3361001.

[ZYQ*08] ZHOU H., YUAN X., QU H., CUI W., CHEN B.: Visual clustering in parallel coordinates. In *Computer Graphics Forum* (2008), vol. *27*, Wiley Online Library, pp. 1047–1054. doi: 10.1111/j.1467-8659.2008.01241.x.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Supporting Information